

(21) Application No 9404888.1

(22) Date of Filing 11.03.1994

(71) Applicant(s)
Motorola Limited

(Incorporated in the United Kingdom)

Jays Close, Viabes Ind Est, BASINGSTOKE,
Hampshire, RG22 4PD, United Kingdom

(72) Inventor(s)
William Neil Robinson

(74) Agent and/or Address for Service
Hugh Christopher Dunlop
Motorola Limited, European Intellectual Property
Operation, Jays Close, Viabes Industrial Estate,
BASINGSTOKE, Hampshire, RG22 4PD,
United Kingdom

(51) INT CL⁶
H04L 1/08 12/28

(52) UK CL (Edition N)
H4P PENL

(56) Documents Cited
None

(58) Field of Search
UK CL (Edition M) H4P PENL PENX PPEC PPG
INT CL⁶ H04L 1/08 1/16 1/18 12/26

(54) Notification by energy burst of messages with unacceptable quality

(57) A data message (Fig. 2 not shown 201) is transmitted 301 to a plurality of communication units (Fig. 1 not shown 103, 105, 107, and 109). The data message (201) is received 401 by the plurality of communication units, which determine 403 whether the quality or anticipated quality of the received data message is acceptable. When at least one of the plurality of communication units determines that the received data message is of unacceptable quality or anticipated quality, the at least one of the plurality of communication units transmits 411 an energy burst (203) in a predetermined time window. Upon detection of the energy burst (203), the transmitting device may retransmit 309 the data message to the plurality of communication units, or may establish a communication link of a selected quality or offer a communication of a different quality.

FIG. 3

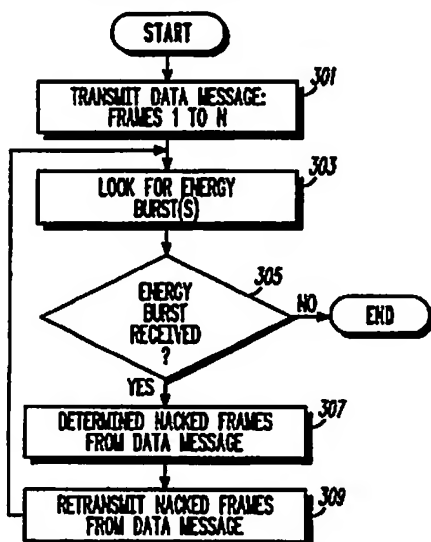
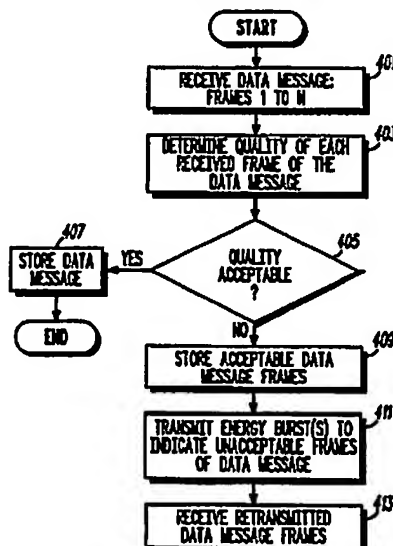


FIG. 4








FIG. 2A



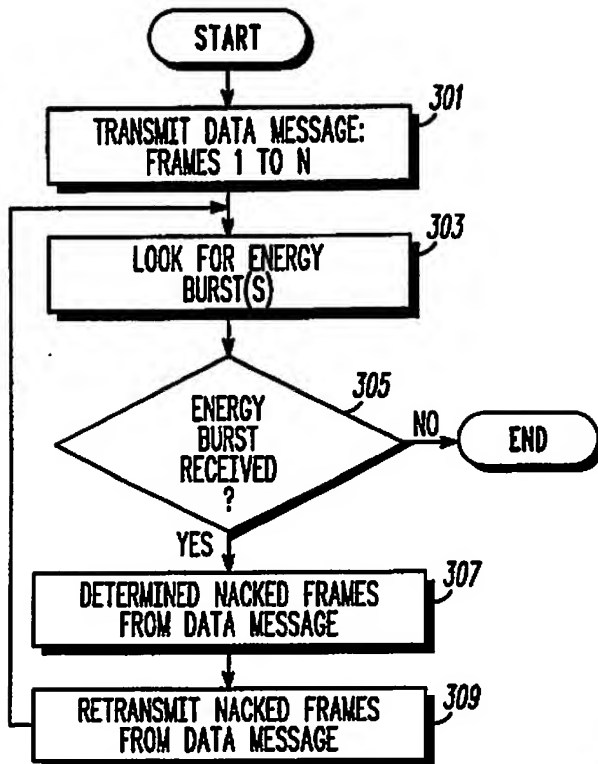
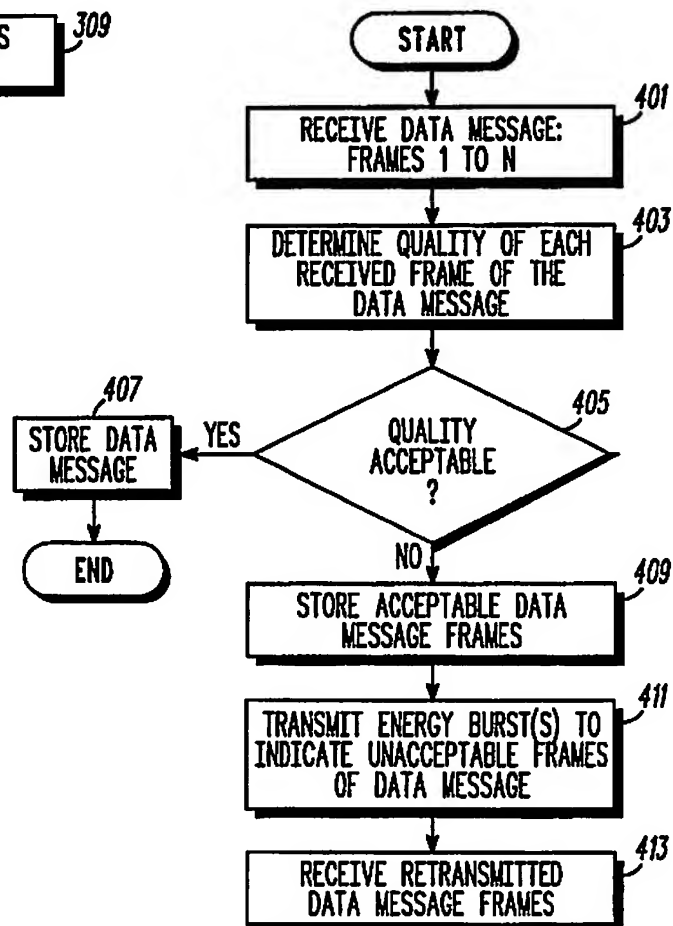



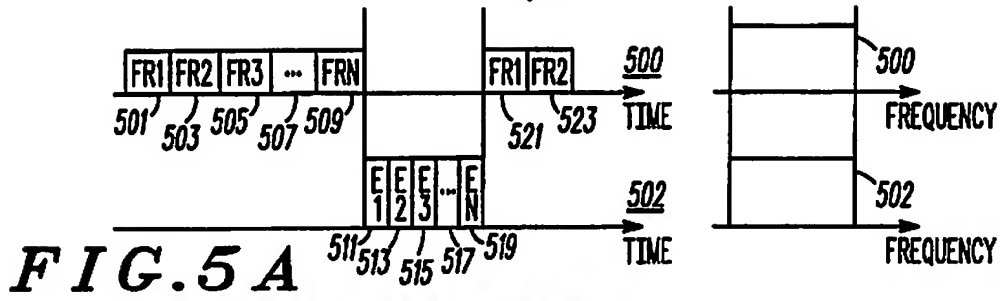


FIG. 2B

FIG. 3**FIG. 4**

3/5



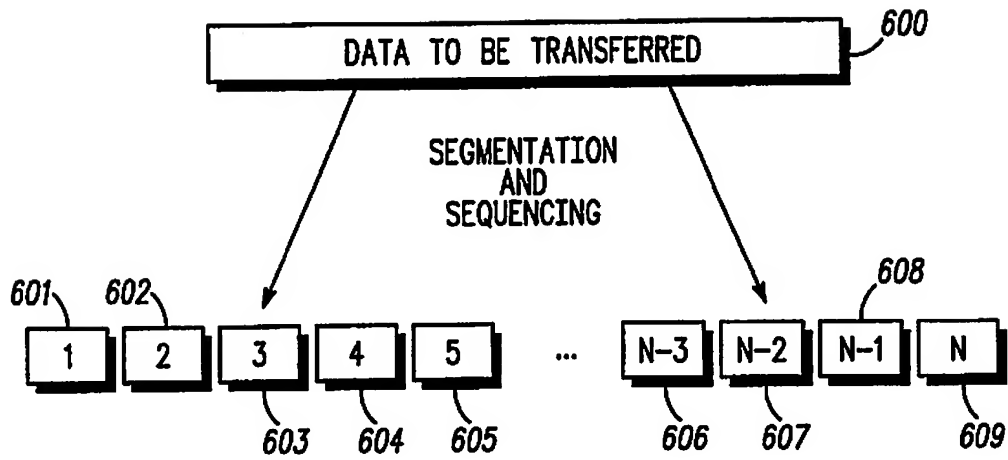


FIG. 6

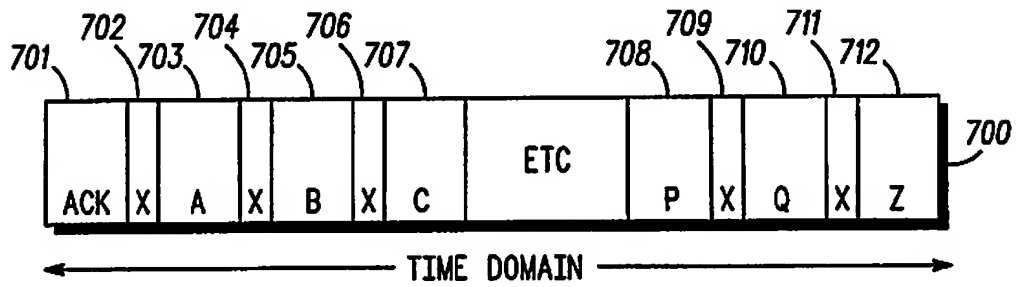


FIG. 7





FIG. 8

FIG. 9

NOTIFICATION BY ENERGY BURST OF UNACCEPTABLE MESSAGES

Field of the Invention

5 This invention relates to data message transmission, including but not limited to notification of actual or anticipated successful receipt of messages.

Background of the Invention

10

 A number of techniques for controlling errors in communication transmissions exist. One such technique is an automatic-repeat-request (ARQ) scheme. The receiver attempts to detect errors according to an error-detection mechanism such as an error-detection code. A data
15 message may be divided into numbered frames. If no errors are detected in a frame of data, the receiver assumes that the frame is error-free, accepts the frame, and via a return channel, may inform the transmitter that the frame was successfully received by identifying the frame by number. A notification of successful reception is often referred to as an ACK, or
20 acknowledgement. If errors are detected in the frame of data, the receiver, via the return channel, may inform the transmitter and request retransmission of the frame by number. This request for retransmission is often referred to as a NACK, or negative acknowledgement. This process of transmission, informing, and retransmission continues until an entire
25 message is successfully received.

 In a broadcast data system (otherwise referred to as a "point to multipoint system"), a single transmission may target numerous receiving communication units, sometimes thousands of units. Because each of the units must send either an ACK or a NACK for each data frame, the return
30 channel becomes inundated with thousands of messages, consuming valuable time on the communication channel. If use of the return channel is not co-ordinated, many of these messages are likely to coincide at least partially, resulting in corrupted ACKs and NACKs, rendering the transmitter unable to distinguish which messages were received
35 successfully and which ones were not. One method for returning ACKs and NACKs via a return channel in a co-ordinated manner is to designate a particular time interval for each unit to transmit an ACK/NACK. When there are numerous communication units, such a system is wasteful of the

communication channel, and messages will take a very long time to be transmitted.

Accordingly, there is a need for an improved method of returning ACKs and NACKs in a broadcast data system utilizing an ARQ protocol.

A connection oriented packet data service can be broadly characterised as having three phases: (1) connect phase; (2) transfer phase and (3) disconnect phase. (A connectionless packet data transfer only has a transfer phase). During the connect phase, the source and destination equipments for the user data sometimes negotiate the establishment of a virtual circuit across which the data will be transferred. As part of this connection establishment, certain defined quality of service (QoS) parameters are sometimes agreed and will be valid for the duration of the virtual circuit. Typical QoS parameters for a connection are widely reported in the literature [e.g. ISO 8348]. Some parameters are appropriate for a point-to-multipoint packet data service. These parameters may be negotiated dynamically or imposed either dynamically or statically by the network or the type of interface.

The negotiation of QoS parameters in a point-to-multipoint system can take a long time or can be inflexible if traditional point-to-point message passing techniques are employed. There is a need for a rapid and reliable manner of completing the connect phase of such a system and there is a need for a flexible manner of managing the quality of service in such a system that is capable of handling mixed technologies and facilitates migration to new technologies.

Summary of the Invention

According to the invention, there is provided a method comprising the steps of: transmitting a data message to a plurality of communication units; receiving the data message by the plurality of communication units; determining, by each of the plurality of communication units, whether the received data message is acceptable; when at least one of the plurality of communication units determines that the received data message is not acceptable, transmitting, by the at least one of the plurality of communication units, an energy burst in a predetermined time window, which time window is the same for each of the plurality of communication units.

A message may be determined as acceptable, for example, if it has

been received with acceptable quality or if it is anticipated that future messages will be received satisfactorily.

The system need not concern itself with which specific units transmit the energy burst (though this can be taken into account if, for example, only a single unit transmits, as is described below). If more than one unit transmits an energy burst, these can be allowed to collide without detriment. The unit transmitting the data message need only determine whether or not an energy burst is transmitted in response to the data message and act accordingly.

One possible action is re-transmission of the data message or part of it. Thus if any one of the plurality of communications units fails to receive the message satisfactorily, it is retransmitted (in whole or in part) to all of them.

Another possible action is a reduction in the quality of service. Thus, in a preferred embodiment of the invention, the data message includes a quality of service indicator indicative of a desired quality of service and the step of determining whether the received data message is acceptable comprises the step of comparing, at each communication unit, the quality of service indicator with a predetermined quality of service indicator for that unit, indicative of the quality of service capability of that unit, and determining that the data message is acceptable if the quality of service of the unit is sufficient for the desired quality of service indicated in the data message.

In this manner, it can be provided that only units not capable of receiving intended data transmit an energy burst. This can be used, for example to negotiate a quality of service for the system. One possible arrangement is for a central unit to offer successively reduced quality of service until no energy burst (NACK) is received. This enables the highest common level of service to be established quickly.

Quality of service may, for example, be data rate, bandwidth, coding type or other type of service.

Since, in each case, the responding communications units respond in the same time window, there is very low overhead usage of the communication channel.

It is not necessary that all communication units in the system are allocated the same time window to respond. In addition to the plurality of communication units, there may be other groups of communication units which are allocated separate time windows.

Other aspects of the invention are defined in the claims. For example, the invention provides, in one aspect, a method of operation of a communications system comprising a central unit and a plurality of remote units, comprising the steps of: transmitting from the central unit a broadcast message including indicators for defining the boolean meaning of a response, receiving the broadcast message at least one remote unit and selectively transmitting from the at least one of the remote units an energy burst in a time window defined by the indicators in the broadcast message. In this way the boolean meaning of information transferred by the selective transmission and non-transmission of an energy burst by the remote unit is dynamically predefined by the broadcast message.

Preferred embodiments of the invention will now be described, by way of example only, with reference to the drawings.

15

Brief Description of the Drawings

FIG. 1 is a block diagram of a communication system in accordance with the invention.

FIG. 2-1 shows timing diagrams and frequency diagrams of the transmission of data messages and energy bursts in the same frequency band in accordance with the invention.

FIG. 2-2 shows timing diagrams and frequency diagrams of the transmission of data messages and energy bursts in different frequency bands in accordance with the invention.

FIG. 3 is a flowchart showing transmitting device activity in accordance with the invention.

FIG. 4 is a flowchart showing receiving device activity in accordance with the invention.

FIG. 5-1 shows timing diagrams and frequency diagrams of the transmission of data messages in frames and transmission of energy bursts in corresponding time windows in accordance with the invention.

FIG. 5-2 shows timing diagrams and frequency diagrams of the transmission of data messages and energy bursts in the same frequency band, and time windows for transmitting energy bursts fall between the transmission of frames of a data message in accordance with the invention.

FIG. 5-3 shows timing diagrams and frequency diagrams of the transmission of data messages, in frames, and energy bursts in different

frequency bands in accordance with the invention.

FIG. 5-4 shows timing diagrams and frequency diagrams of the transmission of data messages, in frames, and energy bursts in separate time windows, where the energy bursts are transmitted in multiple

5 frequency bands in accordance with the invention.

FIG. 6 shows segmentation and sequencing of data into re-transmission units in accordance with a second embodiment of the invention.

FIG. 7 shows a multipoint acknowledge message format in accordance with the second embodiment of the invention.

10 FIG. 8 shows a connect acknowledgement message format in accordance with the second embodiment of the invention.

FIG. 9 shows a typical message sequence in accordance with the second embodiment of the invention.

15 Description of the Preferred Embodiments

The following describes an apparatus for and method of informing a transmitter when a data message was successfully/unsuccessfully received by a plurality of receiving communication units. The communication units
20 that receive a data message determine if the message has acceptable quality. If a unit determines that the message has unacceptable quality, the unit transmits an energy burst in a predetermined time window. When the transmitting unit detects this energy burst, it subsequently retransmits the data message.

25 A data message is transmitted to a plurality of communication units. The data message is received by the plurality of communication units, which determine whether the quality of the received data message is acceptable. When at least one of the plurality of communication units determines that the received data message is of unacceptable quality, at
30 least one of the plurality of communication units transmits an energy burst in a predetermined time window. Upon detection of the energy burst, the transmitting device may retransmit the data message to the plurality of communication units. The predetermined time window may be temporally located immediately subsequent to the receipt of the data message. The
35 energy burst may be a radio frequency energy burst that is transmitted in a predetermined radio frequency band. In addition, the data message may be transmitted in a radio frequency band other than the frequency band in which the energy burst is transmitted, in which instance another message

may be transmitted immediately subsequent to the transmission of the data message.

Alternatively, or in addition, the data message may be comprised of a multiplicity of data frames. The data message is transmitted to a plurality of communication units. The data message is received by the plurality of communication units, which determine whether the quality of each frame of the received data message is acceptable. When at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, at least one communication unit transmits an energy burst in a predetermined time window. The predetermined time window may be temporally located immediately subsequent to the receipt of the data message. The data message may be transmitted in a radio frequency band other than the frequency band in which the energy burst is transmitted. A second data message may be transmitted immediately subsequent to the transmission of the data message.

When the data message is comprised of a multiplicity of data frames, the predetermined time window may be comprised of a multiplicity of time segments. Each frame of the data message may correspond to one of the multiplicity of time segments; and when at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, at least one of the plurality of communication units transmits an energy burst in at least one of the multiplicity of time segments. After detection of the energy burst, the transmitting device may retransmit at least one frame of the received data message to the plurality of communication units. The energy burst may be a radio frequency energy burst that is transmitted in a predetermined radio frequency band.

When the data message is comprised of a multiplicity of data frames, the predetermined radio frequency band may be comprised of a multiplicity of frequency band divisions, wherein each of the multiplicity of data frames corresponds to one of the multiplicity of frequency band divisions; and when at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, at least one of the plurality of communication units transmits an energy burst in at least one of the multiplicity of frequency band divisions corresponding to at least one frame of the received data message.

When the data message is comprised of a multiplicity of data frames,

the predetermined time window may be comprised of a multiplicity of time segments and the predetermined radio frequency band may be comprised of a multiplicity of frequency band divisions, wherein each of the multiplicity of data frames corresponds to a time-frequency pair, wherein each time frequency pair corresponds to one of the multiplicity of time segments and one of the multiplicity of frequency band divisions; and when at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, at least one of the plurality of communication units transmits an energy burst in at least one of the time-frequency pairs corresponding to at least one frame of the received data message.

A communication system is shown in FIG. 1, including a plurality of communication units 103, 105, 107, and 109 and a base station or repeater 101. The communication units, such as a KDT 840 available from Motorola, Inc., are capable of receiving and transmitting data messages.

The base station 101 is shown as having a controller 120 with look-up memory 121 and the mobile stations 105 and 107 are shown (by way of example) as having memory locations 123 and 124. The significance of these is described below with reference to the second embodiment of the invention.

In such a system, a data message 200 is sent from the base station 101 to the communication units 103, 105, 107, and 109. Data messages 200 may also be sourced by a communication unit, a communications console (not shown), or other such device, and sent directly to the other communication units or indirectly via the repeater 101. However the data message 200 is sourced, the data message 200 is transmitted once in a broadcast manner, i.e., a one-to-many transmission, as is known in the art. Only four communication units are shown for simplicity of the drawing, but the present invention will be equally successful in a system of any number of units. Reference numeral 202 represents all energy bursts as transmitted by communication units that do not receive the data message 200 with acceptable quality. Data messages 200 and energy bursts 202 are described in more detail in FIG. 2-1.

A transmitting device may be a base station 101 or a communication unit 103, 105, 107, and 109. A receiving device may also be a base station 101 or a communication unit 103, 105, 107, and 109.

Throughout FIG. 2-1 and FIG. 2-2, the reference numerals 200 or 204 represent any combination of data messages 201 and 205 and the reference

numerals 202 or 206 represent any combination of energy bursts 203 and 207.

Timing diagrams and frequency diagrams of the transmission of data messages 200 and energy bursts 202 are shown in FIG. 2-1, where the data messages and energy bursts are transmitted in the same frequency band. Two data messages 201 and 205 and two energy burst windows 203 and 207 are shown. The timing diagrams on the left of FIG. 2-1 show when data messages and energy bursts are transmitted in relation to one another. For example, a first message 201 is transmitted substantially immediately followed by an energy burst 203, which is transmitted by the receiving communication unit if necessary. A second message 205 may immediately be transmitted upon the end of the time window for the first energy burst 203. A window for a second energy burst 207 follows the completion of the second message 205.

In the frequency diagram on the right of FIG. 2-1, a frequency channel is shown in which both the data messages 200 and the energy bursts 202 occupy the same frequency band width, i.e., the same communication channel.

Timing diagrams and frequency diagrams of the transmission of data messages 204 and energy bursts 206 are shown in FIG. 2-2, where the data messages and energy bursts are transmitted in different frequency bands. Two data messages 201 and 205 may be transmitted back-to-back as shown in the timing diagrams at the left of FIG. 2-2, where, the second message 205 immediately follows the first message 201. The energy bursts 206, if necessary, are transmitted in a predetermined time window following the transmission of the each message, preferably in a time window immediately following or slightly delayed beyond the end of the each message. For example, the first energy burst 203 may be transmitted at the beginning of the transmission of the second message 205. The second energy burst 207 may be transmitted some time immediately after the completion of transmission of the second message 205. The frequency diagram at the right of FIG. 2-2 shows a communication channel that is subdivided, such that data messages 204 are transmitted in the larger section of the communication channel and the energy bursts 206 are transmitted in the smaller section of the communication channel. The frequency band allotted for the energy bursts 206 may be at the beginning (not shown) or the end of the channel, as shown, or even in a different channel all together (not shown), as long as the transmitting units know in

what frequency band to find the energy bursts. For data messages that are comprised of multiple frames and their associated energy bursts, see FIG. 5-1, FIG. 5-2, FIG. 5-3, and FIG. 5-4.

For example, the base station 101 of FIG. 1 transmits two data messages 201 and 205 (200) of FIG. 2 to the communication units 103, 105, 107, and 109. Two communication units 103 and 109 do not receive the first message 201 with acceptable quality and one communication unit 103 does not receive the second message 205 with acceptable quality. The first communication unit 103 transmits two energy bursts 202, an energy burst 203 in the time slot immediately following the first message 201 and an energy burst 207 in the time slot immediately following the second message 205. The second communication unit 109 transmits one energy burst 203 in the time slot immediately following the first message 201. Consequently, both the first communication unit 103 and the second communication unit 109 transmit energy bursts 203 in the first energy burst time window. Because no particular data is included in an energy burst, these two simultaneous transmissions do not interfere with each other in a destructive way but instead reinforce each other, as the base station 101 is merely looking for energy in the time slot, not particular data. Because there is energy in this slot, it is assumed that the previously transmitted message was not successfully received, and the base station knows it was requested to retransmit the message. If all communication units receive a message with acceptable quality, then no energy bursts are transmitted. The base station 101, because no energy is transmitted in the energy burst time window for that message, assumes that the message was received with acceptable quality by all the communication units and will not retransmit the message.

The energy bursts are substantially the same, and the structure of any individual energy burst conveys no message or specific data in itself. Rather, the information conveyed by an energy burst is contained in its temporal location, which indicates the particular data message that was not received with acceptable quality. Thus, if multiple communication units transmit an energy burst simultaneously, they will generally reinforce the effects of other bursts rather than corrupting other bursts. By forcing the simultaneous transmission of energy bursts that serve to NACK a particular data message in accordance with the present invention, a considerable increase in channel efficiency is realized over the prior art.

A flowchart of transmitting device activity is shown in FIG. 3.

Generally, the flow chart of FIG. 3 is embodied in a microprocessor located in the communication units 103, 105, 107, and 109 and/or other transmitting devices, such as base stations 101. A data message comprised of N frames is transmitted at step 301. At step 303, the transmitting device looks for any received energy bursts in a predetermined time window. This window may be located in time immediately following transmission and reception of the data message or delayed some known amount of time after the data message is completely transmitted. Enough time must be allocated to allow communication units at the maximum receivable distance to receive the data message, determine if the message is of acceptable quality, ready the communication unit's transmitter to be set to transmit, and transmit the energy burst. The energy bursts may also be transmitted in a frequency band of the communication channel that is different from the frequency band in which the data message is transmitted. For example, a 500 Hz frequency band may be partitioned from an allocated communication channel for transmission of energy bursts. See FIG. 2-1, FIG. 2-2, FIG. 5-1, FIG. 5-2, FIG. 5-3, and FIG. 5-4 for additional details in the temporal and frequency location of energy bursts with respect to the data messages.

If at step 305 no energy burst is detected, the message is assumed to have been correctly received, and the process ends. If at step 305 one or more energy bursts are detected, then the transmitting device determines which frames were negatively acknowledged (NACKed) at step 307. Because there is a predetermined time and frequency correspondence between each energy burst and a data frame, this determination is simply a correspondence check. The NACKed data frames are retransmitted at step 309, and the process continues with step 303. The transmitting device continues to retransmit NACKed data frames as long as energy bursts are detected in the predetermined time window. The transmitting device may then begin transmission of a new data message. A transmitting device may also transmit many different messages before retransmitting a message, if it is so desired.

A flowchart of receiving device activity is shown in FIG. 4. Generally, the flow chart of FIG. 4 is embodied in a microprocessor located in the communication units 103, 105, 107, and 109 and/or any other receiving devices, such as base stations 101. At step 401, the receiving device receives the N frames of the data message transmitted by the transmitting device in step 301. The quality of each received frame of the data message is determined at step 403. Such determinations are well known in the art,

and include cyclic redundancy checks, Hamming codes, and so forth. If at step 405 the quality of all of the data frames of the data message is acceptable, the process continues with step 407, where the data message is stored in the receiving device, and the process ends. If the quality of one or more of the received data frames is unacceptable at step 405, the process continues with step 407, where the data frames of acceptable quality are stored. At step 411, the receiving device transmits an energy burst for each data frame received with unacceptable quality. The energy burst is transmitted in a predetermined time window, such as shown in FIG. 2-1, FIG. 2-2, FIG. 5-1, FIG. 5-2, FIG. 5-3, and FIG. 5-4. The process continues with step 413, where the retransmitted data message frames are received from step 309, and the process continues with step 403.

The flowcharts of FIG. 3 and FIG. 4 reflect the situation where a data message is divided into frames. The same flowchart applies to a data message that is not divided into frames, such as those shown in FIG. 2-1 and FIG. 2-2, by looking at the unframed message as a framed message having only one frame.

Throughout FIG. 5-1, FIG. 5-2, and FIG. 5-3, the reference numerals 500, 504, 508, or 514 represent all data message frames 501, 503, 505, 507, 509, 521, 523, 525, 527, and 529, and the reference numerals 502, 506, 510, 512, 516, or 518 represent all energy bursts 511, 513, 515, 517, 519, 531, 533, 535, 537, and 539.

Timing diagrams and frequency diagrams of the transmission of data messages and energy bursts are shown in FIG. 5-1, where the data message is transmitted in multiple contiguous frames, and the corresponding energy bursts are transmitted in multiple time segments within a predetermined window of time. Long data messages are commonly divided into smaller sections, typically called frames. As shown in the timing diagrams on the left of FIG. 5-1, the data message is divided into N frames. When transmission of the N frames 501, 503, 505, 507, and 509 is completed, the time window for energy bursts is set substantially immediately after the Nth frame 509. In FIG. 5-1, the time window for the framed message is divided into N time segments, or slots, where each of the N segments corresponds to one of the N frames of the data message. Thus, a time segment is allotted for an energy burst for each of the N frames. One or more of the N energy bursts 511, 513, 515, 517, and 519 are transmitted when the corresponding frame is received with unacceptable quality. A typical data message frame duration is approximately 15 milliseconds,

while a typical energy burst may be approximately 1 millisecond in length. Nevertheless, the energy bursts may be of any length as best suits a particular communication system. After the end of the time window, the transmitter may resume transmitting data message frames 521 and 523 that are part of a new data message or a retransmission of only the particular frame(s) of the first message that were not adequately received.

Timing diagrams and frequency diagrams of the transmission of data messages and energy bursts are shown in FIG. 5-2, where the data messages and energy bursts are transmitted in the same frequency band, and the time windows for transmitting energy bursts fall between the transmission of frames of a data message. The timing diagrams on the left of FIG. 5-2 show an alternative timing pattern to the one shown in FIG. 5-1. A first frame 501 is transmitted following by time slot for a first energy burst 511, a second frame 503 is transmitted followed by a time slot for a second energy burst 513, a third frame 505 is transmitted followed by a time slot for a third energy burst 515, and so forth. Time slots for the data message frames 504 are alternated with the time slots for the energy bursts 506 in this timing scheme. The frequency diagram at the right of FIG. 5-2 shows the message frames 504 and the energy bursts 506 occupying the same frequency band when they are transmitted.

Timing diagrams and frequency diagrams of the transmission of data messages, in frames, and energy bursts in different frequency bands are shown in FIG. 5-3. The frames 508 of the data messages are transmitting back-to-back, i.e., with no time in between, if it is so desired, hence the last frame 509 of a first data message may be followed immediately by the first frame 521 of a second data message. The corresponding energy bursts 510 for each of the frames are transmitted substantially immediately after the frame transmission is completed. For example, the energy burst 511 for the first frame 501 is transmitted substantially immediately after the first frame 501, i.e., during the beginning of the second frame 503. The time segment for the first energy burst 511 may completely fill the time segment allowed for transmission of the second frame 503. Similarly, the energy burst 513 for the second frame 503 may completely fill the time allowed for transmission of the third frame 505. Alternatively, as is seen in the lowest timing diagram, all the energy bursts 512 may be transmitted in sequence following transmission of the Nth frame of the data. Similarly, energy bursts may be sent singly, coupled in pairs or any combination thereof, as long as the transmitting device knows when and where to look for the

energy burst in frequency and time.

The energy bursts 510 and 512 may be transmitted in one or more frequency bands, either within the same communication channel (as shown) or in a different frequency channel (not shown) than the one in which the data message frames 508 are transmitted. As is shown in the frequency diagram on the right of FIG. 5-3, each of the energy bursts is transmitted simultaneously as data is received, but in a different frequency band. The frequency bands for transmitting the energy bursts may be located at the end of the frequency channel (as shown), at the beginning of the channel (not shown), or in the middle of the channel (not shown).

Timing diagrams and frequency diagrams of the transmission of data messages and energy bursts in separate time windows are shown in FIG. 5-4, where the data message is transmitted in multiple contiguous frames, and where the energy bursts are transmitted simultaneously in multiple frequency bands. The timing diagrams on the left of FIG. 5-4 show windows for energy bursts 516 and 518 that are transmitted simultaneously in the same time window. No data frames are transmitted during this time window.

The center frequency diagram on the right of FIG. 5-4 shows that the frequency band in which the data frames 514 are transmitted is divided into N frequency band divisions for energy burst 516 transmission. Each of the N frequency band divisions corresponds to one of the N frames of the data message. An energy burst 516 is simultaneously transmitted in the corresponding one of the N frequency bands for each frame of data that is received with unacceptable quality. At the end of the time window, the transmitter may resume transmitting data message frames 521, 523, 525, and 527, which may be part of a new data message or a retransmission of the particular frame(s) of the first message for which energy bursts were detected.

The lowest timing diagram on the left of FIG. 5-4 shows M energy bursts transmitted simultaneously in each of $N+M$ contiguous time segments, in this case there are two time segments, where M is three, and N is six. The time window is divided into multiple segments, which are shown as being narrow segments, but the duration of the window, and consequently the duration of the time segments, may be increased as necessary.

The lowest frequency diagram on the right of FIG. 5-4 shows that the frequency band in which the data frames 514 are transmitted is divided into M frequency band divisions for energy burst 518 transmission, where M is

three in this example. Each energy burst 518 has both a time location and frequency location, from multiple time locations and multiple frequency locations. Thus, each combination of a single time segment and a frequency band division corresponds to one of the N frames of the data message.

In the example illustrated in FIG. 5-4, the data message frames and the energy bursts are transmitted in substantially the same frequency band. It should be clear to one of ordinary skill in the art that the principles related to use of different frequency bands for the data message and the energy burst transmissions, as illustrated in FIG. 5-3 or FIG. 2-2, may be applied with equal benefit to this case.

Thus, the present invention provides for negative acknowledgement of data messages by the data message receivers that transmit energy bursts, which indicate that part or all of a data message has been received with unacceptably quality. The energy bursts contain no message or specific data. Hence, if multiple communication units transmit an energy burst simultaneously, they will reinforce any other bursts, rather than corrupt or negate a previous energy burst. Because the energy burst is transmitted in a known temporal and spectral location, there is no need to append a particular message number, because the temporal and spectral location of the energy burst identifies the message or message portion. Only a NACK, in the form of an energy burst, is transmitted, hence silence is assumed to be an ACK from all units. Thus, the present invention efficiently uses frequency resources with little waste.

A second embodiment of the invention is now described with reference to FIGS. 6 to 9. For understanding of this embodiment, the following glossary of terms is useful:

BS	base station
CRC	cyclical redundancy check
D-	downlink
FEC	forward error correction
MAT	minimum average throughput
MS	mobile station
QoS	quality of service
RTU	re-transmission unit
Stx	segments transmitted and awaiting acknowledgement
U-	uplink

The second embodiment of the invention utilizes the same equipment

as the first embodiment described above. Referring to FIG. 1, the operation of the system in accordance with the second embodiment of the invention will be described in the example scenario where MS 103 transmits data to the BS 101 on the uplink radio channel. The BS 101 then re-transmits the data on the downlink radio channel, where it is received by MSs 105, 107 and 109. The protocol could equally well be used for direct communications between MS 103 and the other MSs without the intervention of a BS.

For the purposes of explanation, let the term "RTU" be defined as the smallest coherent unit of transmittable data. An RTU can be selectively re-transmitted independently of any other RTU.

If the data is too large to be transmitted in one re-transmission unit (RTU), the data is segmented into a number of RTUs and sequenced as shown in FIG. 6.

FIG. 6 shows a message 600 segmented into n RTUs 601 - 609 numbered 1 to n . Each RTU contains a header plus a segment of data (not shown). The header includes the sequence number of the segment. Other information may be included in the header, according to the needs and capability of the protocol. This other header information is known to those skilled in the art of packet data transfer and is not essential for this explanation.

Each re-transmission unit is transmitted on the downlink. After Stx RTUs have been transmitted on the downlink, where:

$$1 < Stx < Smax$$

and $Smax$ is defined as the maximum number of RTUs that can be transmitted before an acknowledgement is requested from the receiving MSs,

the BS commands the receiving MSs to acknowledge which of the Stx RTUs have been correctly or incorrectly received. This command is called the <D-Acknowledge Request> message.

The <D-Acknowledge Request> message contains the information which identifies some or all of the RTUs that have been transmitted on the downlink and where successful receipt has not previously been acknowledged. Other information may also be included in this message, this other information is known to those skilled in the art and is not essential to this explanation.

The receiving MSs 105-109 are allocated a period of time on the uplink, when the MSs transmit their acknowledgements using a <U-Multipoint Acknowledge> message. The opportunity to transmit the <U-

Multipoint Acknowledge> is simultaneous for all MSs in the cell. For this reason, the design of the <U-Multipoint Acknowledge> message allows the BS 101 to detect for each RTU if one or more MSs requires re-transmission of that RTU. If at least one of the MSs requires the re-transmission of a particular RTU, the BS will re-transmit that RTU.

The design of the <U-Multipoint Acknowledge> message is shown in FIG. 7.

The <U-Multipoint Acknowledge> message comprises the following fields. It has an acknowledge field (Ack) 701 which must be responded to by all MSs which receive the <D-Acknowledge Request>. It has a number of guard times 702, 704, 706, 709 and 711, all labelled x and it has fields 703, 705, 707 (labelled a, b, c) for the first, second and third RTUs in the sequence. Field 708 contains the identification "p" of the (n-1)th RTU in the sequence and field 710 contains the identification "q" of the nth RTU in the sequence. A further field 712 is described below.

The <D-Acknowledge Request> message identifies which RTUs are to be acknowledged by the receiving MSs. The order of the RTUs specified in the <D-Acknowledge Request> message determines the meaning of fields a, b, c p and q in the <U-Multipoint Acknowledge> message. The RTUs identified need not be in order or in sequence. This permits re-transmitted RTUs to be listed amongst RTUs being transmitted for the first time. Thus if the <D-Acknowledge Request> message lists the following RTUs in this order (for example): RTU6; RTU10; RTU11; RTU12; RTU13; RTU15; RTU20; RTU21, then only those RTUs listed can be explicitly acknowledged in the <U-Multipoint Acknowledge> message.

In the corresponding <U-Multipoint Acknowledge> message, the fields take their meaning as defined by the information in the <D-Acknowledge Request> message: field a = RTU6; b = RTU10; c = RTU11, etc. to p = RTU20 and RTU21.

If a receiving MS has not correctly received an RTU identified in the <D-Acknowledge Request> message, then the MS transmits a Negative Acknowledge (NAck) token in the corresponding time period/field. This NAck token can be any appropriate form of radio transmission, for example an analog tone would be quite suitable. If a receiving MS has correctly received the RTU identified in the <D-Acknowledge Request> message, then it shall not transmit anything in fields a to q.

In a further development, the NAck contains an identification of the MS sending the NAck. This has the advantage that, if only one MS sends a

Nack, the identification can be received uncorrupted and stored for later action. One possible action is to ignore that MS, not retransmit any RTUs and communicate the data to that particular MS in some other way (for example at a later, less congested time at a higher power or lower data rate, or by requiring a temporary physical connection to that MS). This has the advantage of not causing repeated retransmissions if a single MS is faulty.

Since the BS may receive NAck tokens from multiple sources during the same time period, it may not be possible to use digital techniques such as FEC and CRC to validate the NAck tokens from each MS. However such techniques are rendered unimportant because the information in the field/time period associated with each RTU in the <U-Multipoint Acknowledge> message only has a binary meaning - NAck or Ack - so the presence or absence of any signal greater than background noise is adequate to determine the binary decision. If the BS receives any increased radio energy during the time periods allocated to the RTU acknowledgements, then it shall re-transmit the associated RTU. As long as at least one MS requires re-transmission of the RTU, then it will be re-transmitted. Those MSs which correctly received the original transmission of the RTU will simply discard the re-transmission.

The Z field 712 in FIG. 7 is used for an MS to indicate using a similar binary token technique that there are other RTUs (not identified in the <D-Acknowledge Request> message which it has not correctly received. The Z field may only be set if the MS has detected that there is an outstanding RTU with an earlier sequence number than the most recent sequence numbers in the <D-Acknowledge Request> message. If the BS detects that the RTUs have not been correctly received, the BS re-transmits the last but one <D-Acknowledge Request> message. If the Z field is again set in the corresponding <U-Multipoint Acknowledge> message, then the BS re-transmits the <D-Acknowledge Request> message which had been sent immediately prior to the last but one <D-Acknowledge Request> message. This continues until the Z field is no longer set. The purpose of the Z field is to overcome the difficulty of MSs not correctly receiving the <D-Acknowledge Request> messages or changing cell.

All MSs which received the <D-Acknowledge Request> message must respond with an Ack token in the Ack field. This Ack token uses the same RF signalling techniques as the NAck field in fields a,b,c, etc. The purpose of the Ack field is to inform the BS whether or not any MSs received the previous <D-Acknowledge Request> message. If no MSS replied with a

<U-Multipoint Acknowledge> message, then this would indicate to the BS that there are no MSs in its coverage area. The BS would then have the choice of repeating the <D-Acknowledge Request> message or terminating the communication in that cell.

5 The guard times 702, 704, 706, 709, 711 are inserted into the <U-Multipoint Acknowledge> message (FIG. 7) to allow the BS to detect the level of background radio interference and to allow for different timing advances arising from the different distances of each MS from the BS. By using different tones for each RTU time period in the <U-Multipoint
10 Acknowledge> message, it is possible to reduce some or all of the guard times to zero, since the BS will be able to detect a change in tone and the BS will be able to detect the transmission activity of the MSs if it samples the uplink signal away from the transition period between adjacent RTU time periods.

15 By reducing the meaning of the RTU acknowledgement fields to a binary value, it is possible to accommodate multiple MSs to simultaneously transmit the same information to the BS using a simple energy token. The meaning of the binary values is assigned by the BS in the previous downlink message. The assignment of the meaning of the binary value is a very
20 useful feature.

 The arrangement described permits point to multipoint packet data to be provided across a radio medium with good reliability and without the need for air interface inefficiencies such as blind re-transmission of data or sequential polling of and individual acknowledgements from each
25 individual receiving MS in the coverage of the BS. Furthermore, the infrastructure does not need to keep track of the location of each MS within the receiving group for paging purposes. There are known techniques for discovering if any MSs are present in the cell without the need for individual paging of MSs, such techniques are known to those skilled in the
30 art and are not essential to this explanation.

 The arrangement described is compatible with many other well known packet data transfer techniques such as "sliding data windows" and "go back N". It can be extended to accommodate any concepts that can be reduced to a binary decision such as "Ready to Receive/Not Ready to
35 Receive".

 It has been shown that reliable and efficient point to multipoint packet data transfer is possible using the <U-Multipoint Acknowledge> method described above. This now makes it possible to achieve the concept

of connection oriented point to multipoint packet data service using the method described above to help maintain the required quality of service.

If one receiving MS in a group moves out of a given coverage area, it will detect more and more errors in the RTUs it receives and will issue
 5 more and more NACK tokens in the <U-Multipoint Acknowledge> message. This in turn will cause more RTU re-transmissions in that cell and will slow down the average transfer rate for all the users in that cell. There will come a point at which the receiving MS in question is requiring so many RTU re-transmissions, that the transfer rate for the other receiving MSs is
 10 too slow. When the receiving MS in question reaches a threshold where it continually requires more than $R_{max} \%$ RTU re-transmissions, it must resign from the connection. This removes the burden of excessive RTU re-transmissions from the BS and increases the achieved average throughput above the maximum average throughput (MAT).

15 There now follows a description of the way in which parameters can be negotiated for a point to multipoint packet data service using aspects of the above described features.

Let R_{max} be any realistic defined function of MAT.

As part of the negotiation of the virtual circuit, the BS transmits a
 20 <D-Connect> message to the receiving MSs using a multipoint (group) address. This message lists the available values for the MAT parameter, plus the preferred value and similarly for other QoS parameters. The parameters may be selected by controller 120 (FIG, 1) from look-up memory
 121.

25 The contents of this message are as follows:

Message identifier : <D-Connect>

Other essential information such as Addresses

MAT selectable value 1 : nn

MAT selectable value 2 : mm

5 MAT selectable value 3 : oo

MAT preferred value : oo

MAT minimum acceptable value : nn

Other QoS parameter 1 selectable value 1 : pp

Other QoS parameter 1 selectable value 2 : qq

10 Other QoS parameter 1 preferred value : qq

Other QoS parameter 1 minimum acceptable value : pp

Other QoS parameter 2 selectable value 1 : rr

Other QoS parameter 2 selectable value 2 : ss

Other QoS parameter 2 preferred value : rr

15 Other QoS parameter 2 minimum acceptable value : ss

etc.

The QoS parameters define, for example, data rate, bandwidth (which may be in terms of frequency bandwidth or number of timeslot in the case of a TDM system), coding rate or coding type, etc.

20 The <D-Connect> message may be repeated for reliability reasons.

The MSs reply with a <U-Connect Ack> message, which is structured in a similar way to the <U-Multipoint Acknowledge> message, but without the Z field (see FIG. 8). As shown in that figure, the fields of the <U-Connect Ack> message are: field 803 = a = RTU6; field 805 = b =
25 RTU10; field 807 = c = RTU11, etc. to field 808 = p = RTU20 and field 810 = q = RTU21.

The meaning of the fields a,b,c etc. are defined according to the order of the available values of the <D-Connect> message. Thus using the definition of the <D-Connect> message above as a template, the a field in the <U-Connect
30 Ack> message would correspond to "MAT selectable value 1", field b would correspond to "MAT selectable value 2", field c would correspond to "MAT selectable value 3". Since "MAT preferred value" and MAT minimum acceptable value" are always one of the MAT selectable values, then the <U-Connect Ack> message does not contain a field which corresponds to these
35 meanings. The receiving MSs decide whether they can accept the preferred values for the negotiable QoS parameters. If this is possible, each receiving MS responds by sending an Ack token in the appropriate time period in the <U-Connect Ack> message. If an MS is unable to comply with one of the

preferred QoS values, it selects a value from the list that it can comply with and sends an Ack token in the appropriate time period in the <U-Connect Ack> message. If a receiving MS is not able to comply with any of the QoS possibilities defined in the <D-Connect> message, then it does not send any

5 Ack tokens and does not participate further in the communication.

Thus, for example, BS controller 120 may initially send a QoS parameter indicative of a transfer rate of 20 RTUs per second. MS 105 has, in memory location 123, a list of transfer rates which that MS is capable of receiving - shown as 20 or 10 in FIG. 1. MS 107, however, only has the

10 transfer rate 10 stored in its memory location 124. Accordingly, MS 107 sends an Ack token with this value in the appropriate time period.

The Ack tokens have the same RF form as the NAck tokens in the <U-Multipoint Acknowledge> message.

On receipt of the Ack token, controller 120 must take a decision as to

15 how to proceed. If the Ack token is received, with the requested transfer rate of 10 RTUs per second, the controller 120 may decide that at least one MS does not have the capability to receive the indicated data rate and may send a new <U-Connect Message> with a reduced rate parameter, e.g. indicative of 10 RTUs per second. If all MSs are able to receive this data

20 rate, none of them will acknowledge in the given time period.

The QoS parameters finally selected for the virtual circuit are chosen after the <U-Connect Ack> response has been received by the BS from the multiple receiving MSs. The values are chosen to give the best overall performance within the limits of the receiving MSs. For example different

25 MSs may have facility for different data rates or different coding types. The objective is to select the best quality of service for which all the MSs have the necessary facility.

These selected QoS parameters may be indicated to the receiving MSs in another downlink message if necessary.

30 The disconnection phase is simply achieved using a <D-Disconnect> message which does not required acknowledgement from the receiving MSs if they have correctly received all of the data. The <D-Disconnect> message may also be repeated for reliability reasons. If an MS receives a

<D-Disconnect> message when it has not correctly received all the data, it can inform the BS by issuing a <U-Disconnect Fail> message, again with the same energy token form as the tokens described earlier. The BS will respond to this by re-transmitting the last <D Acknowledge Request> and
 5 maintaining the virtual circuit. Those MSs which have correctly received the data and the <D-Disconnect> message are permitted to resign from the connection.

It is notable that the point to multipoint connection would continue to exist in a given cell as long as at least one MS continues to receive the data.

10 A typical overall message sequence is shown in FIG. 9

In FIG. 9 a small number of RTUs are transferred from a Base station BS 901 to a multiplicity of mobile stations represented in the figure as MSs 902. These RTUs are transmitted using a point to multipoint connection oriented packet data service. The sequence of message
 15 exchanges begins with the BS 901 broadcasting a <D-Connect> message 903 to the multiplicity of MSs 902. In order to overcome transmissions errors, the <D-Connect> message is repeated as shown in step 904. The BS schedules an opportunity for the multiplicity of MSs 902 to reply simultaneously using the <U-Connect Ack> message 905. Having
 20 considered the <U-Connect Ack> message 905, the BS 901 transmits RTUs 1 to 4 (messages 906, 907, 908 and 909) to the multiplicity of MSs 902. The BS 901 then transmits a <D-Acknowledge Request message 910, inviting the multiplicity of mobiles 902 to acknowledge the correct or incorrect receipt of the previously transmitted RTUs 1 to 4 (messages 906, 907, 908 and 909). The
 25 BS also schedules an opportunity for the multiplicity of mobiles 902 to respond. The multiplicity of mobiles 902 respond with a <U-Multipoint Acknowledge> message 911. In the example shown in FIG 9, one or more of the multiplicity of mobiles 902 indicates in the <U-Connect Acknowledge> 911 message that the transmission of RTU2 907 was unsuccessful. The BS
 30 901 responds to this information by re-transmitting RTU2 912. The BS 901 then issues another <D-Acknowledge Request> message 913 to the multiplicity of mobiles 902 and as before, except that this time only RTU2 (message 912) is indicated as requiring explicit acknowledgement. The multiplicity of mobiles 902 respond at the scheduled time with a <U-
 35 Multipoint Acknowledge> message 914 which in this example indicates that all RTUs have been correctly received by all of the participating mobiles 902. Having learned from the <U-Multipoint Acknowledge> message 914 that all of the RTUs have been correctly received, the BS 901 issues a <D-

Disconnect> message 915 to the multiplicity of mobile 902. This message is repeated 916 to ensure that all of the multiplicity of MSs 902 receive the information.

- 5 Thus the invention provides a point-to-multipoint communication system in which the boolean meaning, i.e. the result of a logical operation, of the information transferred by the transmission or non-transmission of an energy burst is dynamically defined by a previous point to multipoint message according to the rules of the given protocol.

Claims

1. A method comprising the steps of:
 - 5 transmitting a data message to a plurality of communication units;
receiving the data message by the plurality of communication units;
determining, by each of the plurality of communication units, whether the
10 received data message is acceptable;
when at least one of the plurality of communication units determines that
the received data message is not acceptable, transmitting, by the at least
one of the plurality of communication units, an energy burst in a
15 predetermined time window, which time window is the same for each of
the plurality of communication units.
2. The method of claim 1, wherein the step of determining whether the
20 received data message is acceptable comprises the step of checking the
quality of the data message and determining that the data message is
acceptable if the quality thereof is acceptable.
3. The method of claim 2, further comprising the step of retransmitting
25 the data message to the plurality of communication units upon detection of
the energy burst.
4. The method of claim 1, wherein the data message includes a quality
of service indicator indicative of a desired quality of service and the step of
determining whether the received data message is acceptable comprises
30 the step of comparing, at each communication unit, the quality of service
indicator with a predetermined quality of service indicator for that unit,
indicative of the quality of service capability of that unit, and determining
that the data message is acceptable if the quality of service of the unit is
sufficient for the desired quality of service indicated in the data message.
35
5. The method of claim 4, further comprising the step of transmitting a
new data message including a quality of service indicator indicative of a

reduced quality of service to the plurality of communication units upon detection of the energy burst.

6. The method of claim 4 or 5, further comprising the step of
5 transmitting data to the plurality of communication units at a quality of service corresponding to the most recently transmitted quality of service indicator upon detection of no energy burst in the predetermined time window.
- 10 7. The method of any one of claims 3 to 6, wherein the quality of service indicator is indicative of data rate.
8. The method of any one of claims 3 to 6, wherein the quality of service indicator is indicative of bandwidth.
- 15 9. The method of any one of claims 3 to 6, wherein the quality of service indicator is indicative of coding type.
10. The method of any one of the preceding claims, wherein the
20 predetermined time window is temporally located immediately subsequent to the receipt of the data message.
11. The method of any one of the preceding claims, wherein the energy burst includes an identification of the communication unit transmitting the
25 energy burst.
12. The method of claim 11, comprising the step of, upon receiving the energy burst, attempting to identify the communication unit transmitting the energy burst and recording the identity of the communication unit
30 transmitting the energy burst if successfully identified.
13. The method of claim 12, comprising the step of disregarding the identified communication unit and transmitting data to the plurality of communication units at a quality of service corresponding to the most
35 recently transmitted quality of service indicator.
14. The method of claim 13, comprising the further step of later communicating the data to the identified communication unit.

15. The method of any one of the preceding claims wherein the energy burst is transmitted in a predetermined radio frequency band and the data message is transmitted in a radio frequency band other than the frequency band in which the energy burst is transmitted.

16. A method according to claim 1, 2 or 3 wherein the data message comprises a multiplicity of data frames, wherein the step of determining whether the received data message is acceptable comprises determining, by each of the plurality of communication units, whether the quality of each frame of the received data message is acceptable; and

when at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, transmitting, by the at least one of the plurality of communication units, an energy burst in a predetermined time window.

17. The method of claim 16, wherein the predetermined time window is comprised of a multiplicity of time segments.

18. The method of claim 17, wherein each frame of the data message corresponds to one of the multiplicity of time segments; and when at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, transmitting, by the at least one of the plurality of communication units, an energy burst in at least one of the multiplicity of time segments.

19. The method of claim 18, further comprising the step of retransmitting the at least one frame of the received data message to the plurality of communication units after detection of the energy burst.

20. The method of claim 16, wherein the energy burst is a radio frequency energy burst that is transmitted in a predetermined radio frequency band which is comprised of a multiplicity of frequency band divisions, wherein each of the multiplicity of data frames corresponds to one of the multiplicity of frequency band divisions; and when at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, transmitting, by the at

least one of the plurality of communication units, an energy burst in at least one of the multiplicity of frequency band divisions corresponding to the at least one frame of the received data message.

- 5 21. The method of claim 16, wherein the predetermined time window is comprised of a multiplicity of time segments and wherein the predetermined radio frequency band is comprised of a multiplicity of frequency band divisions, wherein each of the multiplicity of data frames corresponds to a time-frequency pair, wherein each time frequency pair
10 corresponds to one of the multiplicity of time segments and one of the multiplicity of frequency band divisions; and when at least one of the plurality of communication units determines that at least one frame of the received data message is of unacceptable quality, transmitting, by the at least one of the plurality of communication units, an energy burst in at
15 least one of the time-frequency pairs corresponding to the at least one frame of the received data message.

22. A communication unit comprising:

- 20 a receiver for receiving a data message;

a determiner, operatively coupled to the receiver, for determining whether the received data message is acceptable; and

- 25 a transmitter, operatively coupled to the determiner, for transmitting an energy burst in a predetermined time window when it is determined that the received data message is not acceptable.

23. A communication unit comprising:

- 30 control means for selecting a desired quality of service and providing at least one desired quality of service indicator,

- a transmitter, operatively coupled to the control means, for transmitting a
35 data message to a plurality of remote communication units, said data message including said at least one quality of service indicator; and

a receiver for detecting an energy burst within a predetermined time

window following transmission of the data message;

wherein the control means are operatively coupled to the receiver and the transmitter to control further reception and transmission of messages at a quality of service which is dependent on reception or non-reception of the energy burst in the predetermined time window.

24. A communication unit according to claim 23, wherein the control means are operatively coupled to the receiver to provide a new quality of service indicator in response to detection of the energy burst, indicative of a different quality of service and for supplying the new quality of service indicator to the transmitter for transmission.

25. A communication unit comprising:

a receiver for receiving a message including a plurality of quality of service indicators,

memory means for storing a predetermined quality of service indicator indicative of a quality of service capability of the communication unit,

control means, operatively coupled to the receiver and the memory means for comparing each of the plurality of quality of service indicators with the predetermined quality of service indicator,

a transmitter, operatively coupled to the control means, for transmitting an energy burst within a predetermined time window dependent on the comparison of the plurality of quality of service indicators and the predetermined quality of service indicator.

26. A communication receiver according to claim 25, wherein the control means are arranged to inhibit the transmitter from transmitting the energy burst and to inhibit further communication when none of the plurality of quality of service indicators matches the predetermined quality of service indicator.

27. A method of operation of a communications system comprising a central unit and a plurality of remote units, comprising the steps of:

- transmitting from the central unit a broadcast message including
indicators for defining the boolean meaning of a response,
receiving the broadcast message at least one remote unit,
selectively transmitting from the at least one of the remote units an
5 energy burst in a time window defined by the indicators in the broadcast
message,
whereby the boolean meaning of information transferred by the
selective transmission and non-transmission of an energy burst by the
remote unit is dynamically predefined by the broadcast message.

Relevant Technical Fields

UK Cl (Ed.M) H4P (PENL, PPEC, PENX, PPG)

(ii) Int Cl (Ed.5) H04L (1/08, 1/16, 1/18, 12/26)

Databases (see below)

(i) UK Patent Office collections of GB, EP, WO and US patent specifications.

(ii)

Search Examiner
 VICKI STRACHAN

Date of completion of Search
 4 MAY 1994

Documents considered relevant
 following a search in respect of
 Claims :-
 1 to 27

Categories of documents

- | | |
|---|---|
| X: Document indicating lack of novelty or of inventive step. | P: Document published on or after the declared priority date but before the filing date of the present application. |
| Y: Document indicating lack of inventive step if combined with one or more other documents of the same category. | E: Patent document published on or after, but with priority date earlier than, the filing date of the present application. |
| A: Document indicating technological background and/or state of the art. | &: Member of the same patent family; corresponding document. |

Category	Identity of document and relevant passages	Relevant to claim(s)
	NONE	

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).